

SupremeRAID™

User Guide for Windows

August 2023



Graid Technology Inc.

Copyright © 2021–2023 Graid Technology Inc. All Rights Reserved.

SupremeRAID™ is a registered trademark of Graid Technology Inc. All other trademarks and registered trademarks are the property of their respective owners.

Graid Technology reserves the right to make changes without further notice to any products or content herein to improve reliability, function, or design. Graid Technology makes no warranty as to the accuracy or completeness of the content or information provided herein, which are provided on an “as is” basis.

No license to Graid Technology’s or any third party’s intellectual property rights are conveyed hereunder.

Publication: August 2, 2023

CHANGE HISTORY

Revision No.	Date	Prepared by /Modified by	Description
1.0.0	2022/12/30	LiamChen	Initial version
1.0.1	2023/8/31	JeffWu	Update user manual for 1.2.3 driver

TABLE OF CONTENTS

CHANGE HISTORY	2
INTRODUCTION	7
Software Module Overview	7
SupremeRAID™ Specifications	7
RAID Components	8
Physical Drive (PD)	8
Drive Group (DG)	8
Virtual Drive (VD)	9
Features Overview	9
Ensuring Data Integrity with Consistency Checks	9
INSTALLATION	10
Prerequisites	10
BIOS Setting Recommend	10
Supported Operating Systems	11
Tested NVMe Devices	11
Installing the Hardware	13
ESD Warning	13
Installation Procedure	13
Installing the Software Driver	15
Upgrading the Software	18
COMMANDS AND SHORTCUTS	19
Syntax	19
Command and Subcommand Quick Reference	20
General	20
Resources	20
Features	22

Managing Licenses	22
Applying the License	22
Checking License Information	23
Checking the SupremeRAID™ Drive Version	24
Checking the Driver Version	24
Resources	24
Viewing Host Drive Information	24
Listing the NVMe Drives	24
Managing Physical Drives	26
Creating a Physical Drive	26
Listing the Physical Drives	27
Deleting the Physical Drives	28
Describing a Physical Drive	29
Locating a Physical Drive	29
Marking a Physical Drive Online or Offline	29
Assigning a Hot Spare Drive	30
Replacing a Nearly Worn-out or Broken SSD	30
Managing Drive Groups	31
Creating a Drive Group	31
Listing Drive Groups	32
Deleting Drive Groups	33
Describing a Drive Group	34
Setting the Drives Group Rebuild Speed	34
Locating Physical Drives in the Drive Group	34
Degradation and Recovery	35
Rescue Mode	35
Managing Virtual Drives	35
Creating a Virtual Drive	35
Listing Virtual Drives	36
Deleting Virtual Drives	38

Adjusting or Updating Configuration Settings for the SupremeRAID™ Add-on	39
Editing Configuration Settings	39
Describing Configuration Settings	39
Deleting Configuration Settings	40
Restoring SupremeRAID™ Configuration Settings	41
Managing Events	42
Listing Events	42
Deleting Events	42
Using Consistency Checks to Ensure Data Integrity	43
Starting Consistency Checks Manually	43
Stopping Consistency Check	45
Scheduling Consistency Checks	45
Viewing Consistency Check Information	46
Setting the Consistency Check Policy	47
Excluding Drive Groups from the Consistency Check Policy	48
ADDITIONAL FUNCTIONS	49
Manually Migrating the RAID Configuration Between Hosts	49
ESXi Virtual Machines Support Using GPU Passthrough	49
Configuring Hosts for NVIDIA GPU Device Passthrough	50
Configuring Virtual Machines	51
Using Self-Encrypting Drives	52
Prerequisites	53
Limitations	53
Importing a Single SED Key Using NQN/WWID	54
Importing a Batched SED Key Using NQN/WWID	54
Displaying SED Key Information	54
Deleting SED Keys	55
TROUBLESHOOTING	56
Sequential Read Performance is not as Expected on a New Drive Group	56

Kernel Log Message "failed to set APST feature (-19)" Appears When Creating Physical Drives.....	56
Installer Error Message: "... nvml.dll not found..." Appears When Installing the SupremeRAID™ Driver	56
Error Message: "graid service is not running"	59
Different LED Blink Patterns on the Backplane.....	62
'graidctl' is not recognized as an internal or external command	63
SPECIFICATIONS	65
SAFETY INFORMATION	66
English Version	66
Chinese Version (SC).....	68
Chinese Version (TC).....	70

INTRODUCTION

SupremeRAID™ is the most powerful, high-speed data protection solution specially designed for NVMe SSDs. SupremeRAID™ installs a virtual NVMe controller onto the operating system and integrates a high-performance, AI processor equipped PCIe RAID card into the system to manage the RAID operations of the virtual NVMe controller.

This document explains how to install the SupremeRAID™ software package for Windows and how to manage the RAID components using the command-line interface.

Software Module Overview

There are two major components of the SupremeRAID™ Software Module:

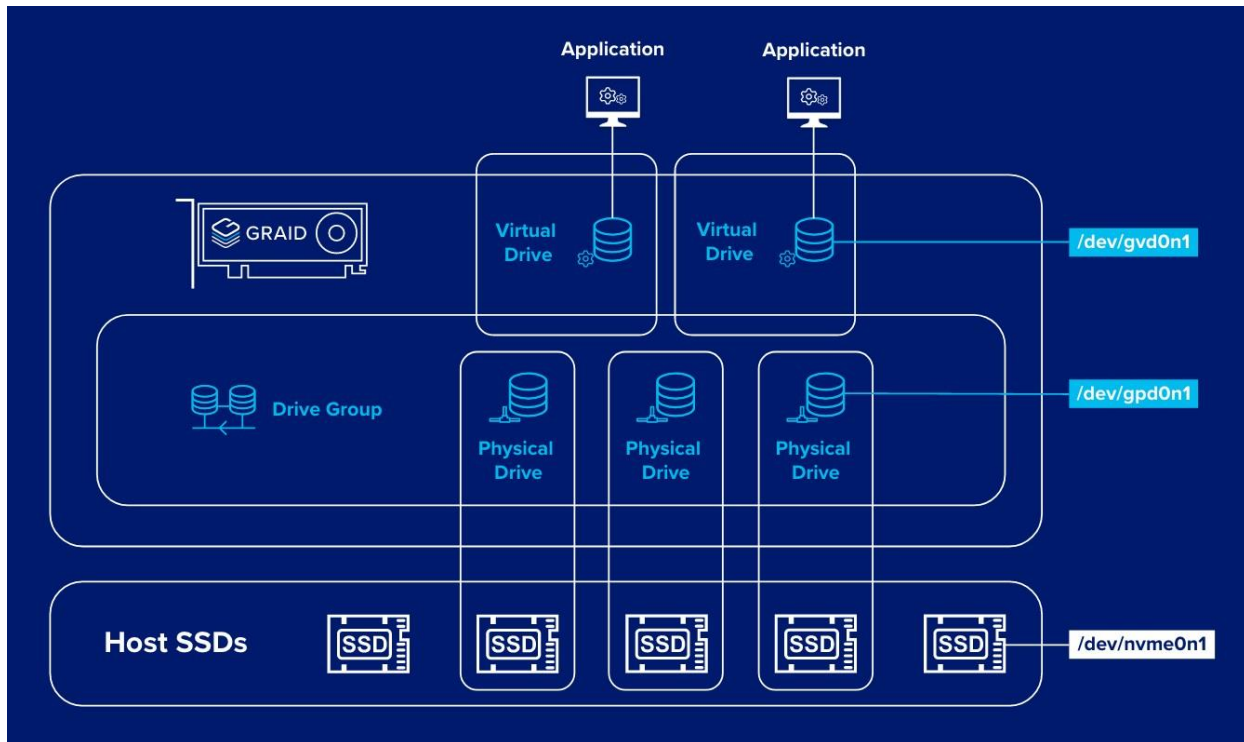
- `graidctl` — The command-line management tool.
- `graid_server` — The management daemon that handles requests from `graidctl` to control the driver.

SupremeRAID™ Specifications

SupremeRAID™ Driver Specifications	
Supported Models	SR-1000, SR-1010, SR-1001
Supported RAID levels	RAID 0, 1, 5, 6, 10,
Recommended minimum drive number for each RAID level	RAID 0 : at least two drives RAID 1 : at least two drives RAID 5 : at least three drives RAID 6 : at least four drives RAID 10 : at least four drives
Maximum number of physical drives	32
Maximum number of drive groups	4
Maximum number of virtual drives per drive group	8
Maximum size of the drive group	Defined by the physical drive sizes

RAID Components

There are three major RAID logical components in SupremeRAID™, the Physical Drive (PD), the Drive Group (DG), and the Virtual Drive (VD).



Physical Drive (PD)

Since NVMe drives are not directly attached to the SupremeRAID™ controller, you must tell the controller which SSDs can be managed. Once an SSD has been created as a physical drive, the SupremeRAID™ driver unbonds the SSD from the operating system, meaning the device node (.\PHYSICALDRIVEx) will disappear and is no longer accessible. At the same time, a corresponding device node is created by the SupremeRAID™ driver. You can check the SSD information, such as the SSD model or SMART logs, using this device node. To control and access the SSD using nvmeXn1, you must first delete the corresponding physical drive.

SupremeRAID™ supports up to 32 NVMe SSD physical drives.

Drive Group (DG)

The main component of RAID logic is a RAID group. When the drive group is created, the SupremeRAID™ driver initializes the physical drives with the corresponding RAID mode to ensure that the data and the parity are synchronized. There are two types of initialization processes.

- **Fast Initialization:** When all the physical drives in the drive group (DG) support the deallocate dataset management command, the SupremeRAID™ driver performs fast initialization by default, meaning the drive group state is optimized immediately.
- **Background Initialization:** Performance will be slightly affected by the initialization traffic, but you can still create the virtual drive and access the virtual drive during a background initialization.

Currently, SupremeRAID™ supports 4 drive groups, with a maximum of 32 physical drives in one drive group.

Virtual Drive (VD)

The virtual drive is equivalent to the RAID volume. You can create multiple virtual drives in the same drive group for multiple applications. The corresponding device node (\\.\PHYSICALDRIVE x) appears on the operating system when you create a virtual drive, and you can make the file system or running application directly on this device node. Currently, the SupremeRAID™ driver supports 8 virtual drives in each drive group.

Features Overview

The SupremeRAID™ presents a range of features that facilitate convenient data storage methods and incorporate diverse protection mechanisms to ensure data integrity. The following will outline key features that contribute to achieving our objectives and fostering a foundational understanding of our services.

Ensuring Data Integrity with Consistency Checks

The SupremeRAID™ is designed to provide high reliability and data integrity levels. A key feature that enables this is the consistency check function.

The consistency check function allows administrators to ensure that the data stored on the SupremeRAID™ system is intact and uncorrupted. These checks can be performed on a regular schedule or manually initiated as needed. When a consistency check is completed, the system compares the data on each disk to identify any discrepancies or errors.

Depending on the settings chosen by the administrator, the consistency check function can either automatically fix any errors that are found or stop the check and alert the administrator to any detected errors. This feature provides administrators with flexibility and control over how the system responds to errors.

For detailed information about graid commands for the consistency check, see [Using Consistency Checks to Ensure Data Integrity](#) on page 44.

Note: The consistency check function is not supported on SupremeRAID™ systems configured in RAID0 mode because RAID0 does not provide data redundancy and does not require data consistency checks.

INSTALLATION

This section describes installing the SupremeRAID™ software package for Windows.

Prerequisites

Before installing the software package, ensure that the system meets the following requirements:

- Minimum system requirements:
 - CPU: 2 GHz or faster with at least 8 cores
 - RAM: 16 GB
 - An available PCIe Gen3 or Gen4 x16 slot
- The SupremeRAID™ card **must** installed into a PCIe x16 slot
- SSD drive

BIOS Setting Recommend

- **Optional** The IOMMU(AMD)/VT-d(intel) function is disabled in the system BIOS(Usually in BIOS Advanced page).
- **Recommended** The UEFI Secure Boot function is disabled in the system BIOS(Usually in BIOS secure page).

Note: The IOMMU(AMD) or VT-d(intel) function must be enabled if you want to use virtualization service like ESXi. Please refer to the [ESXi Virtual Machine Support use GPU Passthrough] section for more information.

Note: Disable UEFI Secure Boot, if not applicable in your system, then you need to be signing the NVIDIA Kernel Module, please refer to the NVIDIA website for additional information and troubleshooting.

Supported Operating Systems

SupremeRAID™ has been tested with the operating system versions in the following table. For other operating system versions, contact Graid Technology support.

Operating System	X64
Windows	10,11*
Windows Server	2019, 2022

***Note:** Windows 11 comes with the "Sleep Mode" enabled by default. However, running the system in "Sleep Mode" might lead to unforeseen errors. To prevent this, it's recommended to manually disable the "Sleep Mode" feature.

Tested NVMe Devices

The following NVMe drives passed Graid Technology Inc. qualification and can be used with SupremeRAID™ at release time. Graid Technology Inc. updates this list when new NVMe drives pass the qualification process.

Manufacturer	Series	Interface	Form Factor
Dapustor	R5100	PCIe Gen4 x 4	2.5 inch U.2
Dapustor	R5103	PCIe Gen4 x 4	2.5 inch U.2
Hagiwara Solution	JN2E	PCIe Gen4 x 4	2.5 inch U.2
Kingston	DC1500M	PCIe Gen3 x 4	2.5 inch U.2
Kioxia	CD5	PCIe Gen3 x 4	2.5 inch U.2
Kioxia	CD6	PCIe Gen4 x 4	2.5 inch U.2
Kioxia	CM6	PCIe Gen4 x 4	2.5 inch U.2
Kioxia	CD8	PCIe Gen4 x 4	2.5 inch U.2
Kioxia	CM7	PCIe Gen5 x 4	2.5 inch U.2
Memblaze	P6536	PCIe Gen4 x 4	2.5 inch U.2

Manufacturer	Series	Interface	Form Factor
Micron	7300 PRO	PCIe Gen4 x 4	2.5 inch U.2
Micron	7450	PCIe Gen4 x 4	2.5 inch U.3
Phison	EPW5900	PCIe Gen4 x 4	2.5 inch U.3
Samsung	PM983	PCIe Gen3 x 4	2.5 inch U.2
Samsung	PM9A3	PCIe Gen4 x 4	2.5 inch U.2
Samsung	PM1733	PCIe Gen4 x 4	2.5 inch U.2
ScaleFlux	CSD-3000	PCIe Gen4 x 4	2.5 inch U.2
Seagate	NYTRO 5550H	PCIe Gen4 x 4	2.5 inch U.2
Solidigm	DC P4510	PCIe Gen3 x 4	2.5 inch U.2
Solidigm	DC P4610	PCIe Gen3 x 4	2.5 inch U.2
Solidigm	D5-P5316	PCIe Gen4 x 4	2.5 inch U.2
Solidigm	D7-P5510	PCIe Gen4 x 4	2.5 inch U.2
Solidigm	D7-P5520	PCIe Gen4 x 4	2.5 inch U.2
Solidigm	D7-P5620	PCIe Gen4 x 4	2.5 inch U.2
Solidigm	Optane™ P5800X	PCIe Gen4 x 4	2.5 inch U.2
Western Digital	SN640	PCIe Gen3 x 4	2.5 inch U.2
Western Digital	SN650	PCIe Gen4 x 4	2.5 inch U.2

For the latest information, see the [Compatible NVMe Drives List](#) on the Graid Technology website.

Installing the Hardware

ESD Warning


Electronic components and circuits are sensitive to Electrostatic Discharge (ESD). When handling any circuit board assemblies including Connect Tech carrier assemblies, it is recommended that ESD safety precautions be observed. ESD safe best practices include, but are not limited to:

- Leaving circuit boards in their antistatic packaging until they are ready to be installed.
- Using a grounded wrist strap when handling circuit boards, at a minimum you should touch a grounded metal object to dissipate any static charge that may be present on you.
- Only handling circuit boards in ESD safe areas, which may include ESD floor and table mats, wrist strap stations and ESD safe lab coats.
- Avoiding handling circuit boards in carpeted areas.
- Try to handle the board by the edges, avoiding contact with components.

Installation Procedure


This section help you to install SupremeRAID™ into your system.

- Step 1** Power down your system.
- Step 2** Unplug the power cord from the AC power source.
- Step 3** Remove the side panel from your system to gain access to the motherboard.
- Step 4** If your system has a PCIe card, remove it. If a retention bar is holding the card in place, remove the screw securing the card. If there is no existing PCIe card, remove the access covers from the primary x16 PCI express slot.



Note: The SupremeRAID™ SR-1010 is dual-slot card and requires you to remove two adjacent slot covers. The SupremeRAID™ SR-1000 and SupremeRAID™ SR-1001 are single slot cards and require only a single-slot.

Step 5 Install the card into the primary x16 PCI Express slot. Press gently on the card until it is seated securely in the slot and reattach the SupremeRAID™ card bracket retention mechanism.



Note: Install the SupremeRAID™ card into the primary x16 PCI Express slot. The SupremeRAID™ SR-1010 is dual-slot card and covers the adjacent slot. The SupremeRAID™ SR-1000 and SupremeRAID™ SR-1001 are single-slot cards. For more information, see <https://manuals.plus/nvidia/rtx-ampere-architecture-based-graphics-card-manual#ixzz7wk7PysLh>.

Step 6 Secure the card to the system frame using the screw(s) you removed in step 4.

Step 7 Install the side panel you removed in step 3.

Installing the Software Driver

Note: You must install the NVIDIA driver before installing the SupremeRAID™ driver for Windows.

To install the SupremeRAID™ driver on Windows systems:

Step 1 Download the latest version of the NVIDIA driver and the SupremeRAID™ driver.

Dependency	
NNVIDIA Driver for Windows 10	537.13
Driver Package	
SR-1001 / SR-1000 / SR1010	Download link

- The SupremeRAID™ software package is downloaded from the Graid Technology, or partner, website.

Note: Please use the appropriate installer based on your specific card model:

SR-1001 : graid-sr-1.2.3-xx.xxxxxxx.001.x64-WHQL.exe

SR-1000 : graid-sr-1.2.3-xx.xxxxxxx.000.x64-WHQL.exe

SR-1010 : graid-sr-1.2.3-xx.xxxxxxx.010.x64-WHQL.exe

Ensure that you download and run the installer corresponding to your respective SupremeRAID™ card installation

Step 2 Install the NVIDIA driver and follow the instructions.

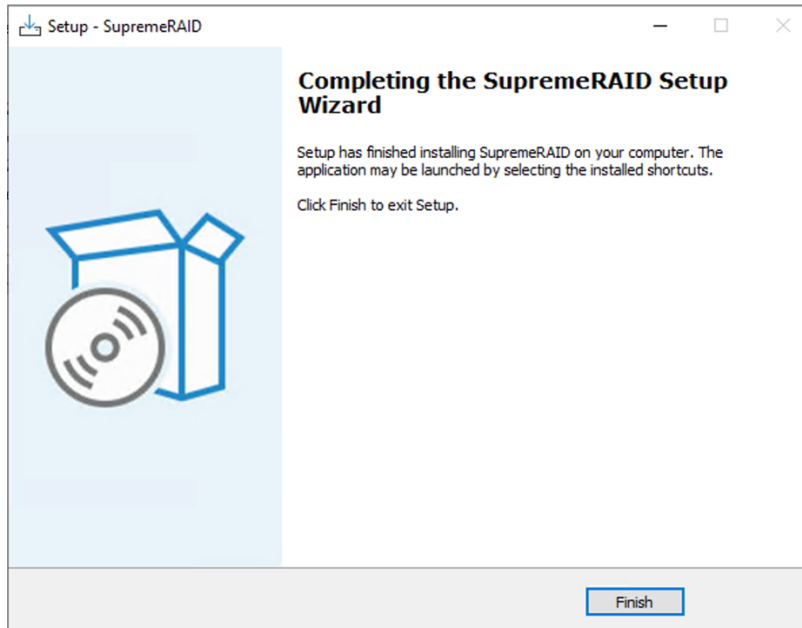


Step 3 Install the SupremeRAID™ driver and follow the instructions.

Step 4 In the end-user license agreement page, you can scroll down the license content. After you review the license, accept the agreement and click Next to proceed.



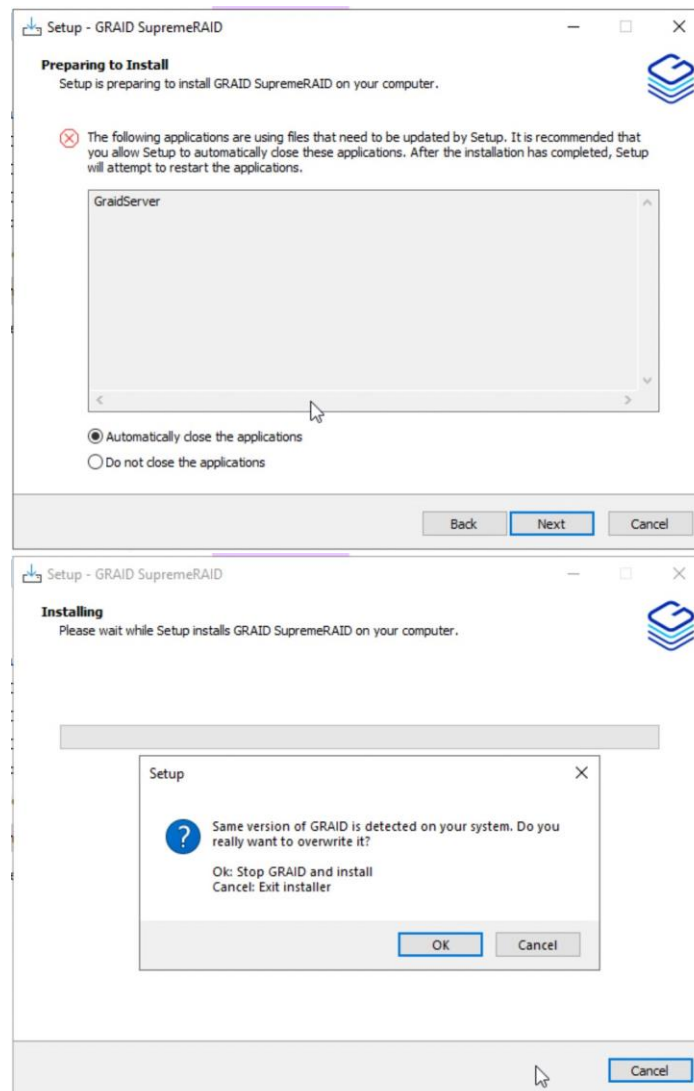
Step 5 Install the SupremeRAID™ driver page.



Upgrading the Software

To upgrade the software:

- Step 1** Download the newest SupremeRAID™ driver and execute the installer.
- Step 2** When the warning message appears, follow the instructions to uninstall the previous driver.



COMMANDS AND SHORTCUTS

Syntax

Use the following syntax to run `graidctl` commands from the terminal window:

```
graidctl [command] [OBJECT_TYPE] [OBJECT_ID] [flags]
```

where `command`, `OBJECT_TYPE`, `OBJECT_ID`, and `flags` are:

- `command`: Specifies the operation to perform on one or more resources, for example create, list, describe, and delete. `OBJECT_TYPE`: Specifies the object type. Object types are case-sensitive, for example `license`, `physical_drive`, and `drive_group`.
- `OBJECT_ID`: Specifies the object ID. Some commands support simultaneous operations on multiple objects. You can specify the `OBJECT_ID` individually, or you can use a dash to describe an `OBJECT_ID` range.

For example, to delete physical drives 1, 3, 4, and 5 simultaneously:

```
> graidctl delete physical_drive 1 3-5
```

- **flags:** Specifies optional flags.

For example:

--force

Force to delete a physical drive

```
> graidctl delete physical_drive 0 --force
```

--json (Supported in coming version)

Print output in json format. This option might assist with API implementation.

```
> graidctl list virtual_drive --format json
```

To get help, run `graidctl help` from the terminal window.

Command and Subcommand Quick Reference

General

Category	Commands	Alias	Sub-Commands	alias
Common	version			
License	apply		license	lic
	describe	desc	license	lic

Resources

Category	Commands	Alias	Sub-Commands	alias
NVMe Drive	list	l, ls	nvme_drive	nd
Physical Drive	create	c, cre, crt	physical_drive	pd
	icreate	ic, icre, icrt	physical_drive	pd
	delete	d, del	physical_drive	pd
	describe	desc	physical_drive	pd

Category	Commands	Alias	Sub-Commands	alias
	edit	e	physical_drive	pd
	list	l, ls	physical_drive	pd
	replace	en	physical_drive	pd
Drive Group	create	c, cre, crt	drive_group	dg
	icreate	ic, icre, icrt	drive_group	dg
	delete	d, del	drive_group	dg
	describe	desc	drive_group	dg
	edit	e	drive_group	dg
	list	l, ls	drive_group	dg
Virtual Drive	create	c, cre, crt	virtual_drive	vd
	icreate	ic, icre, icrt	virtual_drive	vd
	delete	d, del	virtual_drive	vd
	describe	desc	virtual_drive	vd
	edit	e	virtual_drive	vd
	list	l, ls	virtual_drive	vd
Controller	enable		controller	cx
	disable		controller	cx
	delete	d, del	controller	cx
	list	l, ls	controller	cx
	replace	en	controller	cx
Config	describe	desc	config	conf
	edit	e	config	conf
	delete	d, del	config	conf

Category	Commands	Alias	Sub-Commands	alias
	restore	Re	Config	conf
Event	delete	d, del	event	ev
	list	l, ls	event	ev

Features

Category	Commands	Alias	Sub-Commands	alias
Consistency Check	describe	desc	consistency_check	cc
	set		consistency_check	cc
	start		consistency_check	cc
	stop		consistency_check	cc

Managing Licenses

You can apply the license and check license information. To complete the installation, apply the license.

Applying the License

To apply the license and complete the installation, run:

```
> graidctl apply license <LICENSE_KEY>
```

Output example applying invalid license and valid license:

```
C:\Users\Administrator>graidctl apply license 2PE2QH74-UWQXISAH-GHVA48C4-CFBWUSNM
Apply license successfully.
Due to switch GPU to TCC mode, service is inactive now.
Please reboot system to active service.
C:\Users\Administrator>_
```

Note: When applying the license, you might need to provide the serial number of the NVIDIA GPU to Graid Technology Technical Support. To obtain the NVIDIA GPU serial number, run the following command: > `nvidia-smi -q | findstr -i serial`

Note: After applying license, would set NVIDIA driver to TCC mode automatically. You need to reboot the system to enable graid driver.

Checking License Information

To obtain the license information, run:

```
> graidctl describe license
```

Output example:

```
C:\Users\Administrator>graidctl desc license
Describe license successfully.
License State:      APPLIED
License Key:        2PE2QH74-UWQXISAH-GHVA48C4-CFBWUSNM
Expiration Days:    Unlimited
Features:
                    PD Number: 32
                    RAID5: true
                    RAID6: true
                    NVMe-over-Fabric: true
```

Output content:

Field	Description
License State	The current state of your license.
License Key	The applied license key.
Feature	The feature set of the license key.
ExpDays	The expiration date of the license key.

The license state:

State	Description
UNAPPLIED	The license was not applied.
APPLIED	A valid license was applied.
INVALID	A valid license was applied, but a valid RAID card cannot be detected.

Checking the SupremeRAID™ Drive Version

You can prompt the version command to check graidservice information.

Checking the Driver Version

To obtain the graidservice version information, issue:

```
> graidctl version
```

Output example:

```
C:\Users\Administrator>graidctl version
Graidctl version successfully.
graidctl version:      1.2.3-106.gc7fa45ed.010
graid_server version:  1.2.3-106.gc7fa45ed.010
graid kernel driver version: 1.2.3-106
```

Resources

Viewing Host Drive Information

Listing the NVMe Drives

To list all the directly attached NVMe drives, that can be used to create physical drives, run:

```
> graidctl list nvme_drive
```

Output example:


```
C:\Users\Administrator>graidctl list nvme_drive
list nvme drive successfully.
```

DEVICE PATH (4)	MODEL	NQN/WWID	NSID	CAPACITY	ADDRESS
\\?\scsi#disk&ven_nvme&prod_qemu_nvme_ctrl#7&10b6ac7b&0&000000#(53f56307-b6bf-11d0-94f2-00a0c91efb8b)	QEMU NVMe Ctrl	nqn.2019-08.org.qemu:NVME0003	1	1.1 GB	00000:0
\\?\scsi#disk&ven_nvme&prod_qemu_nvme_ctrl#7&28cfbaa2&0&000000#(53f56307-b6bf-11d0-94f2-00a0c91efb8b)	QEMU NVMe Ctrl	nqn.2019-08.org.qemu:NVME0001	1	1.1 GB	00000:0
\\?\scsi#disk&ven_nvme&prod_qemu_nvme_ctrl#7&37ce0091&0&000000#(53f56307-b6bf-11d0-94f2-00a0c91efb8b)	QEMU NVMe Ctrl	nqn.2019-08.org.qemu:NVME0002	1	1.1 GB	00000:0
\\?\scsi#disk&ven_nvme&prod_qemu_nvme_ctrl#7&4766c1&0&000000#(53f56307-b6bf-11d0-94f2-00a0c91efb8b)	QEMU NVMe Ctrl	nqn.2019-08.org.qemu:NVME0004	1	1.1 GB	00000:0

Output content:

Field	Description
License State	The license state.
License Key	The applied license key.
Feature	The feature set of the license key.
ExpDays	The expiration date of the license key.

Managing Physical Drives

Creating a Physical Drive

To create a physical drive, run:

```
> graidctl icreate physical_drive [flag]
```

Or

```
> graidctl c pd < NQN|WWID> [flag]
```

Tip: The "icreate" command enables you select the specific drive to use to create the physical drive.

Note: Only NVMe drives can be used to create physical drives.

Output example for simultaneously creating multiple physical drives with the device path and NQN:

```
C:\Users\Administrator>graidctl icreate physical_drive
? Choose one or more NVMe devices: [Use arrows to move, space to select, <right> to all, <left> to none, type to filter]
> [x] Device Path:\\.\scsi\disk&ven_nvme&prod_qemu_nvme_ctrl#7828cfbaa2&0&0000000#(53f56307-b6bf-11d0-94f2-00a0c91efb8b), NQN/WWID:nqn.2019-08.org.qemu:NVME0001, Model:QEMU NVMe Ctrl, Capacity:1.1 GB
[x] Device Path:\\.\scsi\disk&ven_nvme&prod_qemu_nvme_ctrl#7837ce0091&0&0000000#(53f56307-b6bf-11d0-94f2-00a0c91efb8b), NQN/WWID:nqn.2019-08.org.qemu:NVME0002, Model:QEMU NVMe Ctrl, Capacity:1.1 GB
[x] Device Path:\\.\scsi\disk&ven_nvme&prod_qemu_nvme_ctrl#7810b6ac7b&0&0000000#(53f56307-b6bf-11d0-94f2-00a0c91efb8b), NQN/WWID:nqn.2019-08.org.qemu:NVME0003, Model:QEMU NVMe Ctrl, Capacity:1.1 GB
[x] Device Path:\\.\scsi\disk&ven_nvme&prod_qemu_nvme_ctrl#784766c1&0&0000000#(53f56307-b6bf-11d0-94f2-00a0c91efb8b), NQN/WWID:nqn.2019-08.org.qemu:NVME0004, Model:QEMU NVMe Ctrl, Capacity:1.1 GB

C:\Users\Administrator>graidctl icreate physical_drive
? Choose one or more NVMe devices: Device Path:\\.\scsi\disk&ven_nvme&prod_qemu_nvme_ctrl#7828cfbaa2&0&0000000#(53f56307-b6bf-11d0-94f2-00a0c91efb8b), NQN/WWID:nqn.2019-08.org.qemu:NVME0001, Model:QEMU NVMe Ctrl, Capacity:1.1 GB, Device Path:\\.\scsi\disk&ven_nvme&prod_qemu_nvme_ctrl#7837ce0091&0&0000000#(53f56307-b6bf-11d0-94f2-00a0c91efb8b), NQN/WWID:nqn.2019-08.org.qemu:NVME0002, Model:QEMU NVMe Ctrl, Capacity:1.1 GB, Device Path:\\.\scsi\disk&ven_nvme&prod_qemu_nvme_ctrl#7810b6ac7b&0&0000000#(53f56307-b6bf-11d0-94f2-00a0c91efb8b), NQN/WWID:nqn.2019-08.org.qemu:NVME0003, Model:QEMU NVMe Ctrl, Capacity:1.1 GB, Device Path:\\.\scsi\disk&ven_nvme&prod_qemu_nvme_ctrl#784766c1&0&0000000#(53f56307-b6bf-11d0-94f2-00a0c91efb8b), NQN/WWID:nqn.2019-08.org.qemu:NVME0004, Model:QEMU NVMe Ctrl, Capacity:1.1 GB
? Create new physical drive now? Yes
Icreate physical drive PD0 (nqn.2019-08.org.qemu:NVME0001) successfully.
Icreate physical drive PD1 (nqn.2019-08.org.qemu:NVME0002) successfully.
Icreate physical drive PD2 (nqn.2019-08.org.qemu:NVME0003) successfully.
Icreate physical drive PD3 (nqn.2019-08.org.qemu:NVME0004) successfully.
C:\Users\Administrator>
```

Note: Ensure that the system or other applications are not on the physical drive before creating or replacing it.

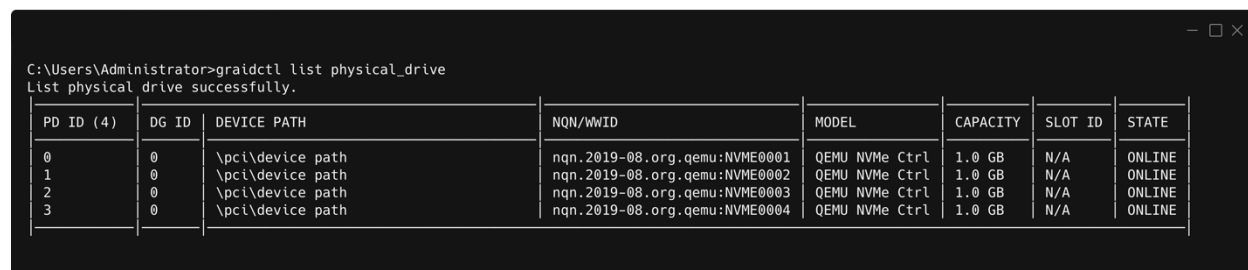
Flag	Description
-h, --help	Help for the list physical_drive command
-f, --dblfwd	Door Bell Forwarding

Listing the Physical Drives

To list all the physical drives, run:

```
> graidctl list physical_drive
```

Output example:



```
C:\Users\Administrator>graidctl list physical_drive
List physical drive successfully.
```

PD ID (4)	DG ID	DEVICE PATH	NQN/WWID	MODEL	CAPACITY	SLOT ID	STATE
0	0	\pci\device path	nqn.2019-08.org.qemu:NVME0001	QEMU NVMe Ctrl	1.0 GB	N/A	ONLINE
1	0	\pci\device path	nqn.2019-08.org.qemu:NVME0002	QEMU NVMe Ctrl	1.0 GB	N/A	ONLINE
2	0	\pci\device path	nqn.2019-08.org.qemu:NVME0003	QEMU NVMe Ctrl	1.0 GB	N/A	ONLINE
3	0	\pci\device path	nqn.2019-08.org.qemu:NVME0004	QEMU NVMe Ctrl	1.0 GB	N/A	ONLINE

Output content:

Field	Description
SLOT ID	This field displays the slot ID of the corresponding NVMe drive. Note that the PD ID is not related to the SLOT ID, and that you must set the physical drives using the PD ID.
DG ID	This field displays the drive group ID of the physical drive.
PD ID	This field displays the PD ID. The PD ID is a unique ID provided by the SupremeRAID™ driver when the physical drive is created. It is not related to any SSD information such as slot ID or NQN. The ID is used for any subsequent operations.
NQN/WWID	This field displays the NQN or WWID of corresponding NVMe drive.
MODEL	This field displays the model number of the corresponding NVMe drive.
CAPACITY	This field displays the capacity of corresponding NVMe drive.
STATE	This field displays the physical drive state.

Physical drive STATE:

State	Description
ONLINE	The physical drive was added to a drive group and is ready to work.
HOTSPARE	The physical drive is configured as hot spare drive.
FAILED	The physical drive is detected but it is not functioning normally.
OFFLINE	The physical drive is marked as offline.
REBUILD	The physical drive is being rebuilt.
MISSING	The physical drive cannot be detected.
INCONSISTENT	The data in the physical drive is inconsistent. This condition generally occurs when the physical drive is in the REBUILD state and the system encounters an abnormal crash.
UNCONFIGURED_GOOD	The physical drive did not join a drive group.
UNCONFIGURED_BAD	The physical drive did not join any drive group and is not functioning normally.

Deleting the Physical Drives

To delete physical drives, run:

```
> graidctl delete physical_drive <PD_ID>
```

Output example for deleting multiple physical drives simultaneously:

```
C:\Users\Administrator>graidctl delete pd 0
Delete physical drive failed: Failed to delete some PDs.
Delete physical drive PD0 failed: PD0 is still used by DG0

C:\Users\Administrator>graidctl delete pd 0-1 3
Delete physical drive successfully.
Delete physical drive PD0 successfully.
Delete physical drive PD1 successfully.
Delete physical drive PD3 successfully.
```

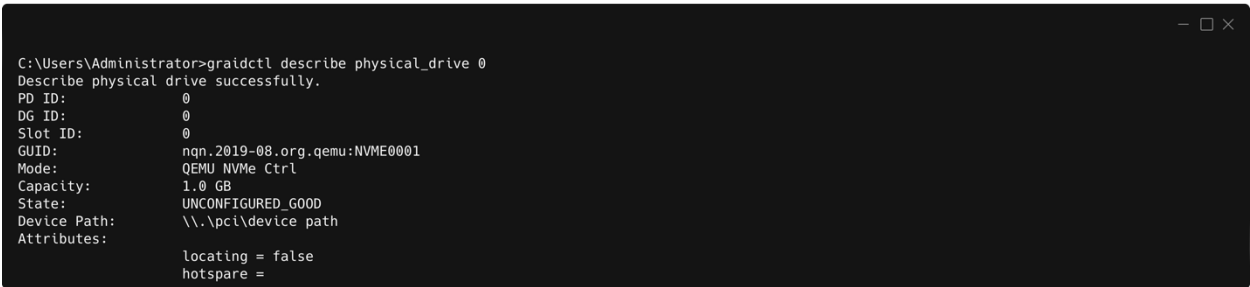
The output shows that a physical drive cannot be deleted when it is part of a drive group.

Describing a Physical Drive

To check the detailed information for a physical drive, run:

```
> graidctl describe physical_drive <PD_ID>
```

Output example:



```
C:\Users\Administrator>graidctl describe physical_drive 0
Describe physical drive successfully.
PD ID:          0
DG ID:          0
Slot ID:        0
GUID:          nqn.2019-08.org.qemu:NVME0001
Mode:          QEMU NVMe Ctrl
Capacity:      1.0 GB
State:         UNCONFIGURED_GOOD
Device Path:   \\.\pci\device path
Attributes:
    locating = false
    hotspare =
```

Locating a Physical Drive

To locate a physical drive, run:

```
> graidctl edit physical_drive <PD_ID> locating start
```

To stop locating a physical drive, run:

```
> graidctl edit physical_drive <PD_ID> locating stop
```

Marking a Physical Drive Online or Offline

To mark a physical drive as online or offline, run:

```
> graidctl edit physical_drive <PD_ID> marker <offline|online>
```

Note: Marking a physical drive as offline, even briefly, puts the physical drive in the REBUILD state.

Assigning a Hot Spare Drive

To assign a physical drive as global hot spare, run:

```
> graidctl edit physical_drive <PD_ID> hotspare global
```

To assign a physical drive as the hot spare for a specific drive group, run:

```
> graidctl edit physical_drive <PD_ID> hotspare <DG_ID>
```

To assign a physical drive as a hot spare for multiple drive groups, use a comma (,) to separate the drive group IDs.

Replacing a Nearly Worn-out or Broken SSD

Step 1 Mark the physical drive as bad using the following command. (You can skip this step if the physical drive is in the MISSING or other abnormal state.)

```
> graidctl edit pd <OLD_PD_ID> marker bad
```

Step 2 Replace the NVMe SSD. The state of the prior physical drive will indicate FAILED.

Step 3 Check the NQN of the new SSD.

```
> graidctl list nvme_drive
```

Step 4 Create a new physical drive on the new SSD.

```
> graidctl create physical_drive <NEW_SSD_NQN>
```

Step 5 Replace the physical drive.

```
> graidctl replace physical_drive <OLD_PD_ID> <NEW_PD_ID>
```

Step 6 Delete the old physical drive.

```
> graidctl delete physical_drive <OLD_PD_ID>
```

Managing Drive Groups

Creating a Drive Group

To create a drive group, run:

```
> graidctl create drive_group <RAID_MODE> (PD_IDs) [--background-init]
```

Output example showing three drive groups were created:

```
PS C:\> graidctl create drive_group raid1 0-1
Create drive group DG0 successfully.
PS C:\> graidctl create drive_group raid5 2-4
Create drive group DG1 successfully.
PS C:\> graidctl create drive_group raid6 5-9
Create drive group DG2 successfully.
```

Required parameters:

Option	Description
RAID_MODE	This field specifies the RAID mode of the drive group. Entries must be all uppercase or all lowercase. (For example, RAID6 or raid6 are both correct)
PD_IDs	This field specifies the ID of the physical drive joining the drive group.

Optional parameters:

Option	Description
backgroundinit	Default option. Use standard methods to initialize the drive group. When all the physical drives in the drive group support the deallocate dataset management command, it is used to synchronize the data, or parity, between the physical drives during the creation of the drive group.

Note: Wait for the drive group initialization to complete. DO NOT power-off or reboot the system when the drive_group state is INIT/RESYNC/RECOVERY.

Use the command below to check drive_group state:

```
> graidctl list drive_group
```

Listing Drive Groups

To list all drive groups, run:

```
> graidctl list drive_group
```

Output example:

```
PS C:\> graidctl list drive_group
List drive group successfully.
```

DG ID	MODE	VD NUM	CAPACITY	FREE	USED	STATE
0	RAID1	0	3.2 TB	3.2 TB	0 B	OPTIMAL
1	RAID5	0	6.4 TB	6.4 TB	0 B	OPTIMAL
2	RAID6	0	9.6 TB	9.6 TB	0 B	OPTIMAL

Output content:

Field	Description
DG ID	This field displays the drive group ID.
MODE	This field displays the drive group RAID mode.
VD NUM	This field displays the number of virtual drives in the drive group.
CAPACITY	This field displays the total usable capacity of the drive group.
FREE	This field displays the unused space of the drive group.
USED	This field displays the used space of the drive group.
STATE	This field displays the drive group state.

Drive Group STATE:

STATE	Description
OFFLINE	The drive group does not function normally. This condition is usually caused when the number of damaged physical drives exceeds the limit.
OPTIMAL	The drive group is in optimal state.
DEGRADED	The drive group is available and ready, but the number of missing or failed physical drives has reached the limit.
PARTIALLY_DEGRADED	The drive group is available and ready for use, but some physical drives are missing or failed.
RECOVERY	The drive group is recovering.
FAILED	The drive group does not function normally.
INIT	The drive group is initializing.
RESYNC	The drive group is re-synchronizing. This condition usually occurs when the system encounters an abnormal crash. Do not replace the physical drive in this state until the resynchronization process is complete.
RESCUE	The drive group is in rescue mode.

Deleting Drive Groups

To delete drive groups, run:

```
> graidctl delete drive_group <DG_ID>
```

```
C:\Users\Administrator>graidctl delete drive_group 2
Delete drive group failed: Failed to delete some DGs.
Delete drive group DG2 failed: DG0 still has IVD(s)

C:\Users\Administrator>graidctl delete dg 0-1 3
Delete drive group successfully.
Delete drive group DG0 successfully.
Delete drive group DG1 successfully.
Delete drive group DG3 successfully.
```

You cannot delete a drive group that contains a virtual drive.


In this example, drive group 1 was not deleted because it contains a virtual drive. Drive groups 0 and 2 were deleted successfully.

Describing a Drive Group

To display detailed information for a drive group, run:

```
> graidctl describe drive_group <DG_ID>
```

Output example:



```
C:\Users\Administrator>graidctl describe drive_group 0
Describe drive group successfully.
DG ID:          0
Mode:           RAID5
Capacity:       3.0 GB (3019898880 B)
Free Space:     3.0 GB (3019898880 B)
Used Space:     0 B
State:          OPTIMAL
PD IDs:         [0 1 2 3]
Number of VDs:  0
Attributes:     rebuild_speed = high
```

Setting the Drives Group Rebuild Speed

To set the rebuild speed for a drive group, run:

```
> graidctl edit drive_group <DG_ID> rebuild_speed {low|normal|high}
```

Locating Physical Drives in the Drive Group

To locate all the physical drives in the drive group, run:

```
> graidctl edit drive_group <DG_ID> locating start
```

To stop locating all the physical drives in drive group, run:

```
> graidctl edit drive_group <DG_ID> locating stop
```

Degradation and Recovery

When multiple drive groups require simultaneous recovery, the drive groups recover individually.

When multiple physical drives in the same drive group require rebuilding, the physical drives are rebuilt simultaneously.

Rescue Mode

When a damaged drive group is initialized, or when a recovering drive group encounters an abnormal system crash, the data integrity of the drive group is affected. In this event, the drive group is forced offline to prevent data from being written to the drive group. To read the data for the drive group, force the drive group to go online using Rescue mode.

Note: A drive group in Rescue mode is read-only. Rescue mode cannot be disabled.

To enter the rescue mode, run:

```
> graidctl edit drive_group <DG_ID> rescue_mode on
```

Managing Virtual Drives

Creating a Virtual Drive

To create a virtual drive, run:

```
> graidctl create virtual_drive <DG_ID> [<VD_SIZE>]
```

Output example:

```
PS C:\> graidctl create virtual_drive 0
Create virtual drive VD0 in DG0 successfully.
PS C:\> graidctl create virtual_drive 1 100G
Create virtual drive VD0 in DG1 successfully.
PS C:\> graidctl create virtual_drive 2 1T
Create virtual drive VD0 in DG2 successfully.
```

Listing Virtual Drives

To list virtual drives, run:

```
> graidctl list virtual_drive [--dg-id=<DG_ID>] [--vd-id=<VD_ID>]
```

Output example:

```
PS C:\> graidctl list virtual_drive
List virtual drive successfully.
```

VD ID (4)	DG ID	SIZE	DEVICE PATH	STATE
0	0	3.2 TB	\\.\PHYSICAKDRIVE1	OPTIMAL
0	1	100 GB	\\.\PHYSICAKDRIVE2	OPTIMAL
0	2	1.0 TB	\\.\PHYSICAKDRIVE3	OPTIMAL
1	2	1.0 TB	\\.\PHYSICAKDRIVE4	OPTIMAL

Output content:

Field	Description
DG ID	This field displays the drive group ID.
VD ID	This field displays the virtual drive ID.
SIZE	This field displays the usable size of the virtual drive.
DEVICE PATH	This field displays the device path of the virtual drive.
NQN	This field displays the NQN of the virtual drive.
STATE	This field displays the virtual drive state. It is identical to the drive group state.
EXPORTED	This field displays whether the virtual drive was exported using NVMe-oF or iSCSI.

Note: Do not perform I/O before the virtual drive is initialized and the device path (for example, /dev/gvd0n) is created.

Virtual Drive STATE: Identical to the drive group state.

STATE	Description
OFFLINE	The drive group does not function normally. This condition is usually caused when the number of damaged physical drives exceeds the limit.
OPTIMAL	The drive group is in optimal state.
DEGRADED	The drive group is available and ready, but the number of missing or failed physical drives has reached the limit.
PARTIALLY_DEGRADED	The drive group is available and ready for use, but some physical drives are missing or failed.
RECOVERY	The drive group is recovering.
FAILED	The drive group does not function normally.
INIT	The drive group is initializing.
RESYNC	The drive group is re-synchronizing. This condition usually occurs when the system encounters an abnormal crash. Do not replace the physical drive in this state until the re-synchronization process is complete.
RESCUE	The drive group is in rescue mode.

Deleting Virtual Drives

To delete virtual drives, run:

```
> graidctl delete virtual_drive <DG_ID> <VD_ID> [--force]
```

Output example:

```
PS C:\> graidctl delete virtual_drive 0 0
Delete virtual drive VD0 from DG0 successfully.
PS C:\> graidctl delete virtual_drive 2 0-1
Delete virtual drive VD1 from DG2 successfully.
Delete virtual drive VD0 from DG2 successfully.
```

The example shows that a virtual drive being used by the application cannot be deleted without adding the force flag.

To create a RAID-5 virtual drive with 5 NVMe SSDs:

Step 1 Create a physical drive.

```
> graidctl icreate physical_drive
```

```
Create physical drive PD0 successfully.
Create physical drive PD1 successfully.
Create physical drive PD2 successfully.
Create physical drive PD3 successfully.
Create physical drive PD4 successfully.
```

Step 2 Create a drive group.

```
> graidctl create drive_group raid5 0-4
Create drive group DG0 successfully.
```

Step 3 Create a virtual drive.

```
> graidctl create virtual_drive 0
```

Create virtual drive VD0 successfully.

Step 4 Check the device path of the new virtual drive.

```
> graidctl list virtual_drive --dg_id=0
```

		DG ID	VD ID	SIZE	DEVICE PATH	STATE
		0	0	500 GB	\\.\PHYSICALDRIVE1	OPTIMAL

Adjusting or Updating Configuration Settings for the SupremeRAID™ Add-on

The add-on for SupremeRAID™ provides enhanced configuration options and allows you to fine-tune system settings to meet your specific needs. Follow these steps to ensure that the add-on is configured optimally for maximum system performance.

Editing Configuration Settings

To edit the configuration, issue the following command:

```
> graidctl edit config <config_name> <value> [flags]
```

or

```
> graidctl e conf <config_name> <value> [flags]
```

Flag	Description
-h, --help	Help for the edit config command

Configuration Options

Field	Description
SED	Add single SED key for specific device

Output Example

```
C:\Users\Administrator>graidctl edit config sed nqn.2019-08.org.qemu:NVME0001
Enter Key: Edit config successfully.
```

Describing Configuration Settings

To describe the configuration, issue the following command:

```
> describe config <config_name> [flags]
```

OR

```
> desc conf <config_name> [flags]
```

Flag	Description
-h, --help	Help for the describe config command

Configuration Options

Field	Description
SED	Obtain the SED key information

Output Example

```
C:\Users\Administrator>graidctl describe config sed
Describe config successfully.
Totally 1 SED keys.
Device GUIDs:
    nqn.2019-08.org.qemu:NVME0001
```

Deleting Configuration Settings

To delete the configuration, issue the following command:

```
> graidctl delete config <config_name> [flags]
```

OR

```
> graidctl del conf <config_name> [flags]
```

Flag	Description
-h, --help	Help for the delete config command

Configuration Options

Field	Description
SED	Obtain the SED key information

Output Example

```
[graid@graid demo~]$ sudo graidctl delete config sed all
Do you really want to delete all SED key?
Repeat IMEANTODELETEALL to continue: IMEANTODELETEALL
✓Delete config successfully.
[graid@graid demo~]$ sudo graidctl del conf led
✓Delete config successfully.
```

Restoring SupremeRAID™ Configuration Settings

To scan all NVMe and SCSI drives and restore the latest SupremeRAID™ configuration, issue the following command:

```
> graidctl restore config [flags]
```

OR

```
> graidctl re conf [flags]
```

Flag	Description
-h, --help	Help for the restore config command
-a, --auto	Selects the last configuration automatically

Output Example

```
[graid@graid demo~]$ sudo graidctl restore config
✖Restore config failed: Please stop the graid service before restoring the config, and restart the graid service after restored the config.
[graid@graid demo~]$ sudo graidctl re conf
Skip /dev/sda: no config found
Found the following configs:
0: Device /dev/nvme0n1, UUID 00200000-0000-0000-4d02-000000000000, Epoch 1412, Time 2022-12-08 20:14:09 +0800 CST
1: Device /dev/nvme1n1, UUID 00200000-0000-0000-4d02-000000000000, Epoch 1412, Time 2022-12-08 20:14:09 +0800 CST
2: Device /dev/nvme2n1, UUID 00200000-0000-0000-4d02-000000000000, Epoch 1412, Time 2022-12-08 20:14:09 +0800 CST
3: Device /dev/nvme3n1, UUID 00200000-0000-0000-4d02-000000000000, Epoch 1412, Time 2022-12-08 20:14:09 +0800 CST
Please select one config to restore (0-3): 0
Restore to /etc/graid.conf (y/N)? y
✓Restore config graid.conf successfully.
```

Managing Events

Listing Events

To check detailed information from record, issue the following command:

```
> graidctl list event [flags]
```

OR

```
> graidctl ls event [flags]
```

Flag	Description
-h, --help	Help for the list event command
-c, --component	[string] Filter events by component
-n, --max_entries	[int32] Limit the number of events returned
-o, --output	[string] Output to a file
-s, --severity	[string] Filter events by severity

Output Example

```
C:\Users\Administrator>graidctl list event -n 10 -s INFO -c dg
List event successfully.
[2023-08-31 23:00:36 -0700 PDT][INFO][DG][2] State transitted from UNKNOWN to OPTIMAL.
[2023-08-31 23:00:36 -0700 PDT][INFO][DG][2] Drive group created.
[2023-08-31 23:00:47 -0700 PDT][INFO][DG][3] State transitted from UNKNOWN to OPTIMAL.
[2023-08-31 23:00:47 -0700 PDT][INFO][DG][3] Drive group created.
[2023-08-31 23:01:05 -0700 PDT][INFO][DG][0] Drive group deleted.
[2023-08-31 23:01:05 -0700 PDT][INFO][DG][1] Drive group deleted.
[2023-08-31 23:01:05 -0700 PDT][INFO][DG][3] Drive group deleted.
[2023-08-31 23:04:04 -0700 PDT][INFO][DG][2] Drive group deleted.
[2023-08-31 23:04:18 -0700 PDT][INFO][DG][0] State transitted from UNKNOWN to OPTIMAL.
[2023-08-31 23:04:18 -0700 PDT][INFO][DG][0] Drive group created.
```

Deleting Events

To delete events, issue the following command:

```
> graidctl delete event [flags]
```

OR

```
> graidctl del event [flags]
```

Flag	Description
-h, --help	Help for the delete event command
-d, --date	[string] Delete event entries before the date
-e, --entries	int32] Keep the latest number of entries Default: -1

Using Consistency Checks to Ensure Data Integrity

The consistency check operation verifies that the data is correct in DGs that use RAID levels 1, 5, 6, and 10. In a system with parity, for example, checking consistency calculates the data on one drive and compares the results to the contents of the parity drive.

Note: You cannot perform a consistency check on RAID 0 because it does not provide data redundancy. Additionally, a consistency check can only run when the DG is in OPTIMAL or PARTIALLY_DEGRADED state.

The consistency check function records all events to the event database, and graidctl provides commands to retrieve the events. The maximum number of event entries is 1,000. The system deletes event entries periodically. You can also delete entries manually.

Starting Consistency Checks Manually

To start a consistency, check manually, issue the following command:

```
> graidctl start consistency_check manual_task [flags]
```

OR

```
> graidctl start cc [flags]
```

Flag	Description
-h, --help	Help for the start consistency_check manual command
-p, --policy	[string] Specify CC policy [stop_on_error/auto_fix]

DG State

Enabling a consistency check task adds the following annotations beside the output string of the DG state.

DG State	Description
OPTIMAL	Normal state without enabling consistency check
OPTIMAL (!)	Inconsistency found
OPTIMAL (cc)	Consistency check ongoing
OPTIMAL (cc!)	Consistency check ongoing and inconsistency found

Output Example

```
C:\Users\Administrator>graidctl start consistency_check manual_task 0 1 -p stop_on_error
Start consistency check successfully.

C:\Users\Administrator>graidctl start cc manual_task 0 -p auto_fix
Start consistency check successfully.
```

Stopping Consistency Check

To stop a consistency check task, issue the following command:

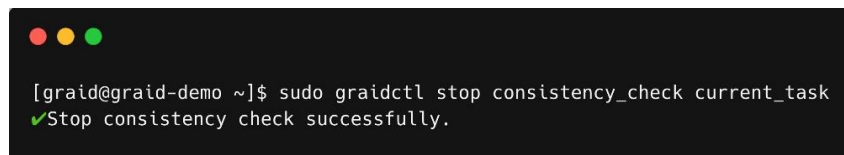
```
> graidctl stop consistency_check current_task [flags]
```

or

```
> graidctl stop cc current_task [flags]
```

Flag	Description
-h, --help	Help for the stop consistency_check command

Output Example



```
[graid@graid-demo ~]$ sudo graidctl stop consistency_check current_task
✓Stop consistency check successfully.
```

Scheduling Consistency Checks

To schedule a consistency check task, issue the following command:

```
> graidctl set consistency_check schedule_mode
<off|continuously|hourly|daily|weekly|monthly><yyyy/mm/dd> <hh> [flags]
```

or

```
> graidctl set cc schedule_mode
<off|continuously|hourly|daily|weekly|monthly> <yyyy/mm/dd> <hh> [flags]
```

Flag	Description
-h, --help	Help for the set consistency_check command

DG State

Enabling a consistency check task adds the following annotations beside the output string of the DG state.

DG State	Description
OPTIMAL	Normal state without enabling consistency check
OPTIMAL (!)	Inconsistency found
OPTIMAL (cc)	Consistency checks ongoing
OPTIMAL (cc!)	Consistency checks ongoing and inconsistency found

Output Example

```
C:\Users\Administrator>graidctl set cc schedule_mode daily 2023/08/29 10
Set consistency check successfully.
```

Viewing Consistency Check Information

To view detailed consistency check information, issue the following command:

```
> graidctl describe consistency_check [flags]
```

or

```
> graidctl desc consistency_check [flags]
```

Flag	Description
-h, --help	Help for the describe consistency_check command

Output Example

```
C:\Users\Administrator>graidctl desc consistency_check current_task
Describe consistency check successfully.
Schedule Mode:      daily
Schedule Base:      2023-08-29 10:00:00 -0700 PDT
Excluded DGs:       []
Policy:              stop_on_error
Next Schedule:      2023-08-29 10:00:00 -0700 PDT
```

Setting the Consistency Check Policy

To set a consistency check policy, issue the following command.

Note: By default, the consistency check runs on all drive_groups. To exclude drive groups, run the `xcluded_dgs` command.

```
> graidctl set consistency_check policy <auto_fix|stop_on_error> [flags]
```

Flag	Description
-h, --help	Help for the set consistency_check command

Output Example

```
C:\Users\Administrator>graidctl set consistency_check policy auto_fix
Set consistency check successfully.
```

Excluding Drive Groups from the Consistency Check Policy

To exclude some drive groups from a consistency check policy, issue the following command:

```
> graidctl set consistency_check excluded_dgs <DG_IDs>
```

or

```
> graidctl set cc excluded_dgs <DG_IDs>
```

Flag	Description
-h, --help	Help for the set consistency_check command

Output Example

```
C:\Users\Administrator>graidctl set consistency_check excluded_dgs 0
Set consistency check successfully.
```


ADDITIONAL FUNCTIONS

Manually Migrating the RAID Configuration Between Hosts

To manually migrate the RAID configuration between hosts:

Step 1 Periodically backup the configuration file %PROGRAMFILES%\graid SupremeRAID\conf\graid.conf from the original host. To move the configuration file to another system.

Step 2 Setup the target host and ensure that the graid service is stopped.

Note: When the target host already contains an installed and running SupremeRAID™ card, stop and restart the service using the graid.conf file from the original system.

Step 3 Move all SSDs from the original host to the new host.

Step 4 Copy the configuration backup file to the new host using the same path.

Step 5 Start the graid service directly if the original card also moved to the new host.

```
> sc start graidservice
```

Otherwise, you must apply the new license of the new SupremeRAID™ card on the new host.

```
> graidctl apply license <LICENSE_KEY>
```


ESXi Virtual Machines Support Using GPU Passthrough

You can create virtual machines with SupremeRAID™ support to maximize performance.

Hypervisor support: VMWare ESXi 7.0U3 / ESXi8.0

Configuring Hosts for NVIDIA GPU Device Passthrough

Step 1 Put ESXi host into maintenance mode. From the Navigator menu, select Host > Enter maintenance mode.

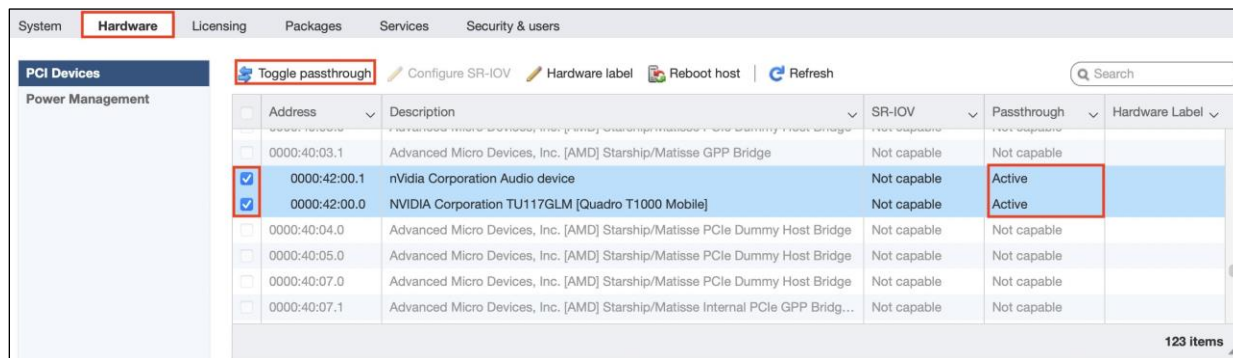


Step 2 Manage the PCI device passthrough. From the Navigator menu, select Manage > Hardware > PCI Devices. The Passthrough Configuration page appears listing all of the available pass-through devices.

Step 3 Select the NVIDIA T1000 (Quadro T1000 Mobile) and its audio device.

Step 4 Click Toggle passthrough.


Step 5 Check the Passthrough status. It should be Active.



Configuring Virtual Machines

Step 1 Attach PCI devices to virtual machines. From the Edit VM setting page, choose Virtual Hardware > Add other device > PCI device.

Step 2 Choose Quadro T1000 and its Audio device as the two PCI devices.

▶  PCI device 1	TU117GLM [Quadro T1000 Mobile] - 0000:42:00.0	✕
▶  PCI device 2	<class> Audio device - 0000:42:00.1	✕


Note: When the T1000 PCI device is assigned to the virtual machine, you must set the memory reservation to accommodate the fully configured memory size.


Step 3 Choose Virtual Hardware > Memory.


Step 4 Check Reserve all guest memory (All locked).




Virtual Hardware

VM Options

 Add hard disk

 Add network adapter

 Add other device

▶  CPU	8	▼	
▼  Memory			
RAM	16	GB	▼
Reservation	16384	▼	MB ▼
<input checked="" type="checkbox"/> Reserve all guest memory (All locked)			

Step 5 Enable point-to-Point (P2P) on the Virtual Machine for best performance. From the Edit VM setting page, choose VM Options > Advanced > Configuration Parameters > Edit Configuration....

Virtual Hardware VM Options

General Options VM Name:

VMware Remote Console Options ☐ Lock the guest operating system when the last remote user disconnects

VMware Tools Expand for VMware Tools settings

Power management Expand for power management settings

Boot Options Expand for boot options

Advanced

Settings ☐ Disable acceleration ☒ Enable logging

Debugging and statistics Run normally

Swap file location ☒ Default Use the settings of the cluster or host containing the virtual machine. ☐ Virtual machine directory Store the swap file in the same directory as the virtual machine. ☐ Datastore specified by host Store the swap files in the datastore specified by the host to be used for swap files. If not possible, store the swap files in the same directory as the virtual machine. Using a datastore that is not visible to both hosts during vMotion might affect the vMotion performance for the affected virtual machines.

Configuration Parameters

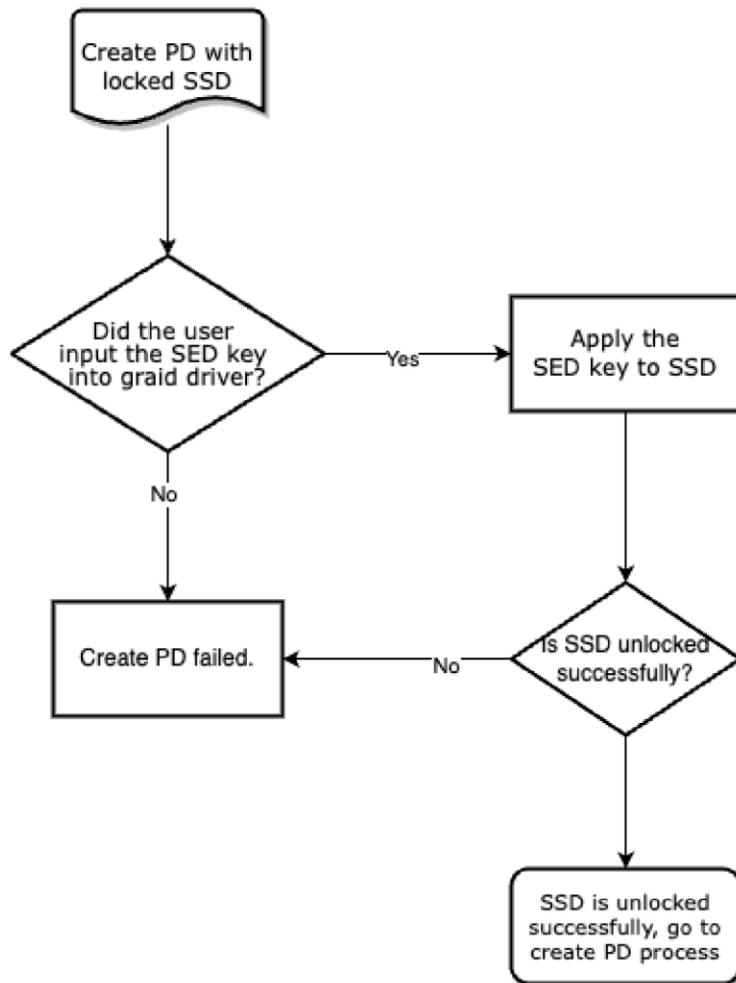
Step 6 Add below parameters.

```
hypervisor.cpuid.v0 = FALSE
pciPassthru.allowP2P = TRUE
pciPassthru.use64bitMMIO = TRUE
pciPassthru.64bitMMIOSizeGB = 64
```

Using Self-Encrypting Drives

A self-encrypting drive (SED) uses native full-disk encryption. SupremeRAID™ supports SEDs and SED key management. When the SED key is configured, SupremeRAID™ uses the imported key to unlock the SED.

Note: You must configure the SED key using the graidctl tool before creating the physical drives.



Prerequisites

- Collect the NQN/WWID of the NVMe disks.(They are required to import the SED key.)
- Prepare the SED key for each disk.

Limitations

- Only NVMe devices are supported.
- Only the "global" range parameter is supported.

Importing a Single SED Key Using NQN/WWID

To import a single SED key using NQN/WWID, issue the following command:

```
> graidctl edit config sed <NQN/WWID>
```

The following figure shows an example.



```
graid@graid:~$ sudo graidctl edit config sed nqn.2014-08.org.nvmexpress:uuid:52bbdb40-c5bf-f92d-9961-a6368e845bfd
Enter Key:
✓ Edit config sed successfully.
```

Importing a Batched SED Key Using NQN/WWID

To import a batched SED key using NQN/WWID, issue the following command:

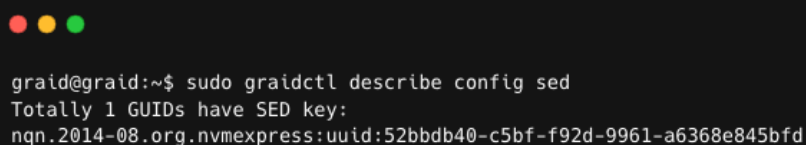
```
> graidctl edit config sed file <filename> file content format:
<NQN1/WWID1>, <KEY1>
<NQN1/WWID1>, <KEY2>
...
<NQNm/WWIDm>, <KEYm>
```

Displaying SED Key Information

To display SED key information, issue the following command:

```
> graidctl describe config sed
```

The following figure shows an example.



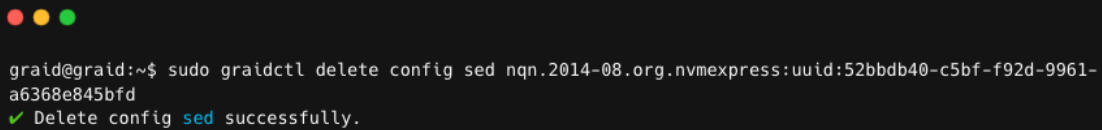
```
graid@graid:~$ sudo graidctl describe config sed
Totally 1 GUIDs have SED key:
nqn.2014-08.org.nvmexpress:uuid:52bbdb40-c5bf-f92d-9961-a6368e845bfd
```

Deleting SED Keys

To delete a SED key, issue the following command:

```
> graidctl delete config sed <GUID>
```

The following figure shows an example.

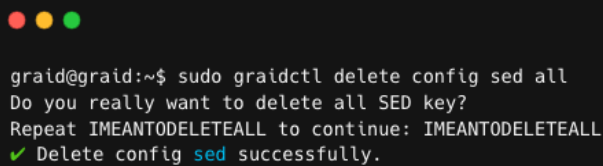


```
graid@graid:~$ sudo graidctl delete config sed nqn.2014-08.org.nvmexpress:uuid:52bbdb40-c5bf-f92d-9961-a6368e845bfd
✓ Delete config sed successfully.
```

To delete all SED keys, issue the following command:

```
> graidctl delete config sed all
```

The following figure shows an example.



```
graid@graid:~$ sudo graidctl delete config sed all
Do you really want to delete all SED key?
Repeat IMEANTODELETEALL to continue: IMEANTODELETEALL
✓ Delete config sed successfully.
```

TROUBLESHOOTING

Sequential Read Performance is not as Expected on a New Drive Group

Unlike SAS/SATA hard drives, many NVMe SSDs support the deallocate dataset management command. Using this command, you can reset all data in the NVMe SSD immediately, eliminating the need to synchronize data between physical drives when creating a drive group.

But for other SSDs, the performance is not as expected when reading unwritten sectors after issuing the deallocate dataset management command. While this behavior also impacts the performance of the new drive group, it does not affect the applications because they do not read sectors that do not contain data.

To test SupremeRAID™ performance, write the entire virtual drive sequentially using a large block size.

Kernel Log Message "failed to set APST feature (-19)" Appears When Creating Physical Drives


Some NVMe SSD models might display a "failed to set APST feature (-19)" message in the kernel log when creating the physical drive.

When SupremeRAID™ creates the physical drive, the SSD is unbound from the operating system so that SupremeRAID™ can control the SSD. During the unbinding process, when the APST feature is enabled, the NVMe driver attempts and fails to set the APST state to SSD, and the error message is issued.

This message is expected and can be ignored. SupremeRAID™ is functioning normally.

Installer Error Message: "... nvml.dll not found..." Appears When Installing the SupremeRAID™ Driver

Error message:

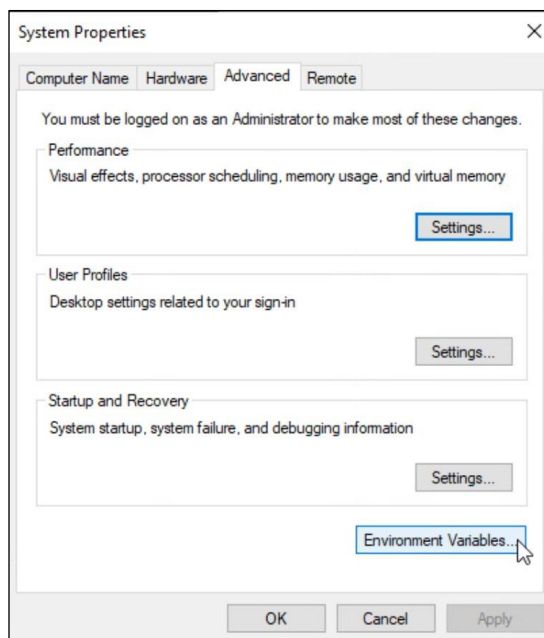


During installation the SupremeRAID™ driver must detect the presence of the NVIDIA driver. However, the NVIDIA driver does not automatically append the installation folder to the environment variables during its installation. This condition prevents the SupremeRAID™ driver from detecting the NVIDIA driver and causes the error message. Appending the NVSMI folder path to the environment variables fixes the issue.

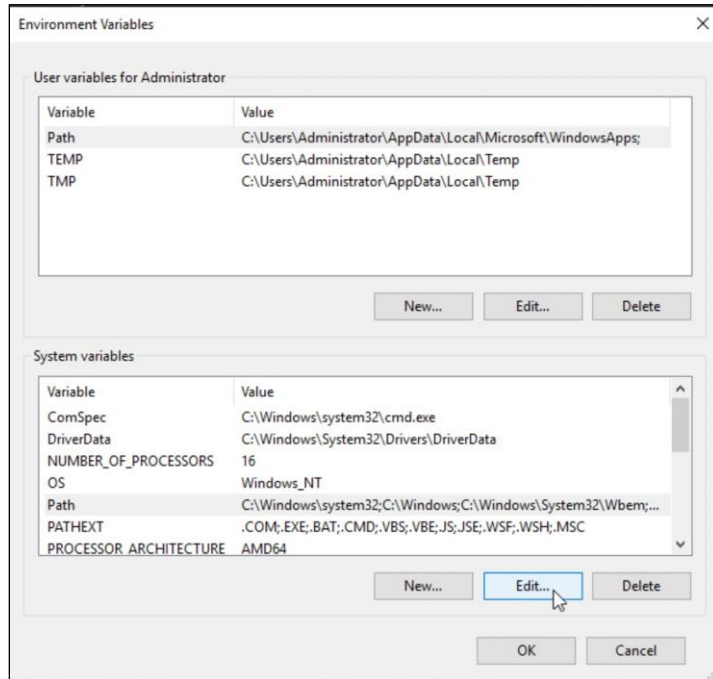
To append the NVSMI folder path to the environment variables:

Step 1 Press "Windows Key + R" and type sysdm.cpl to open the System Properties dialog.

Step 2 Select the Advanced tab and click Environment Variables....

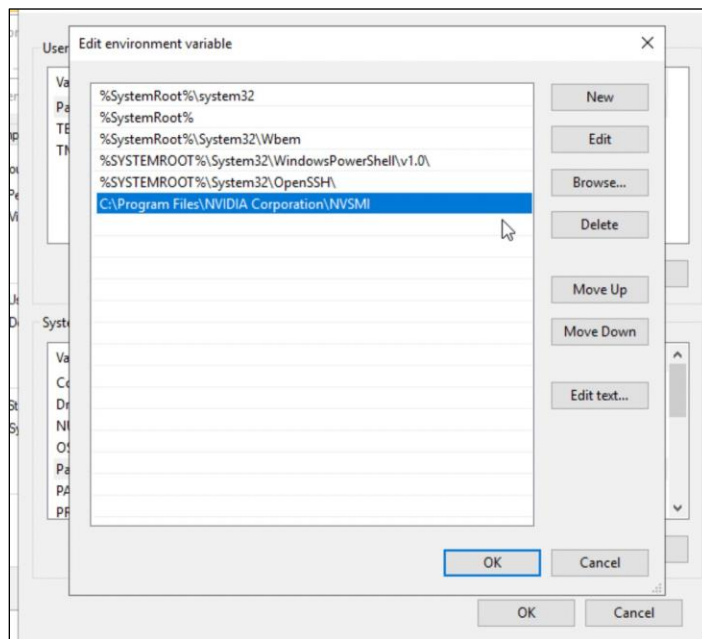


Step 3 From the Environment Variables dialog, select the Path row and click Edit.



Step 4 Add the NVSMI path, save it, and reboot the system.

NVSMI path: %PROGRAMFILES%\NVIDIA Corporation\NVSMI\



Error Message: "graid service is not running"

Use either of the following methods to check the status of the graid service.

To check the graid service using a command prompt or PowerShell:

Step 1 From a command prompt or PowerShell, run:

```
sc.exe query graidserver
```

Step 2 If the service is not "RUNNING", type the following command and wait a moment before checking its status again.

```
sc.exe start graidserver
```

Output example:

```
PS C:> graidctl list nvme_drive
GRAID service is not running

PS C:> sc.exe query graidserver

SERVICE_NAME: graidserver
        TYPE               : 10  WIN32_OWN_PROCESS
        STATE                : 1  STOPPED
        WIN32_EXIT_CODE       : 0  (0x0)
        SERVICE_EXIT_CODE    : 0  (0x0)
        CHECKPOINT           : 0x0
        WAIT_HINT            : 0x0

PS C:> sc.exe start graidserver

SERVICE_NAME: graidserver
        TYPE               : 10  WIN32_OWN_PROCESS
        STATE                : 2  START_PENDING
                           (NOT_STOPPABLE, NOT_PAUSABLE, IGNORES_SHUTDOWN)
        WIN32_EXIT_CODE       : 0  (0x0)
        SERVICE_EXIT_CODE    : 0  (0x0)
        CHECKPOINT           : 0x1
        WAIT_HINT            : 0x1d4c0
        PID                 : 4500
        FLAGS                 :

PS C:> sc.exe query graidserver

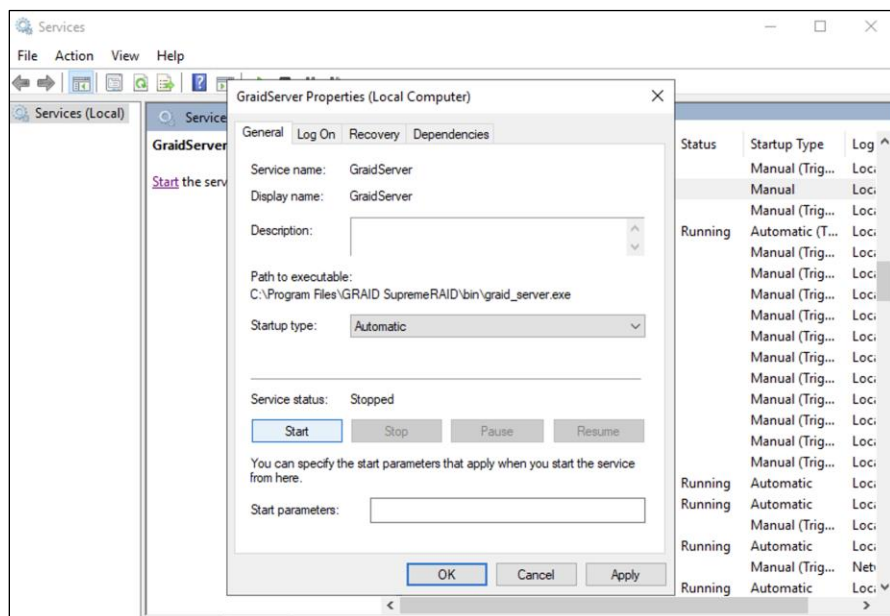
SERVICE_NAME: graidserver
        TYPE               : 10  WIN32_OWN_PROCESS
        STATE                : 4  RUNNING
                           (STOPPABLE, NOT_PAUSABLE, ACCEPTS_PRESHUTDOWN)
        WIN32_EXIT_CODE       : 0  (0x0)
        SERVICE_EXIT_CODE    : 0  (0x0)
        CHECKPOINT           : 0x0
        WAIT_HINT            : 0x0
```

To check the graid service using Windows Service Manager

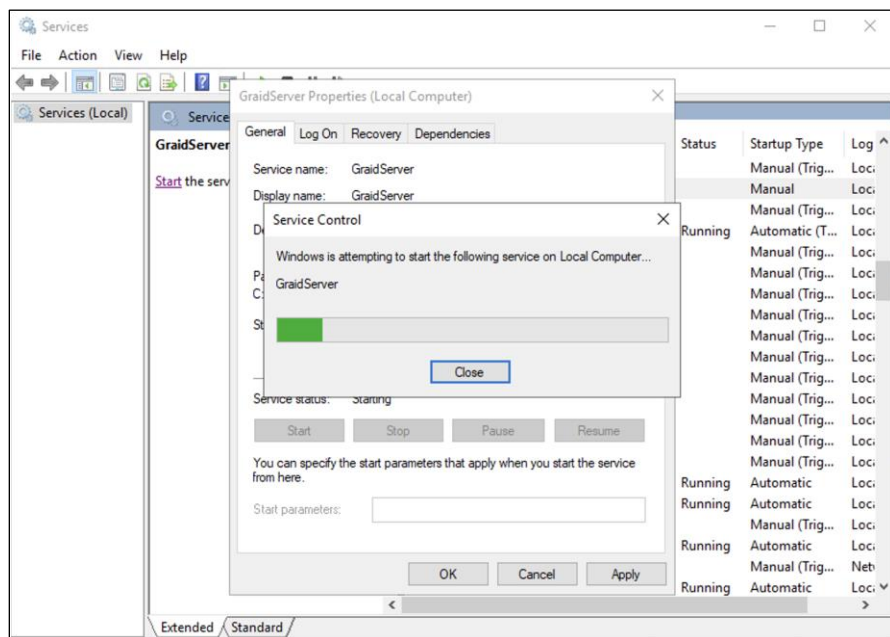
Step 1 Open Windows Service Manager and search "graidserver" to check its status.

Step 2 If the service is not "RUNNING", click Start and wait a moment for the process to complete.

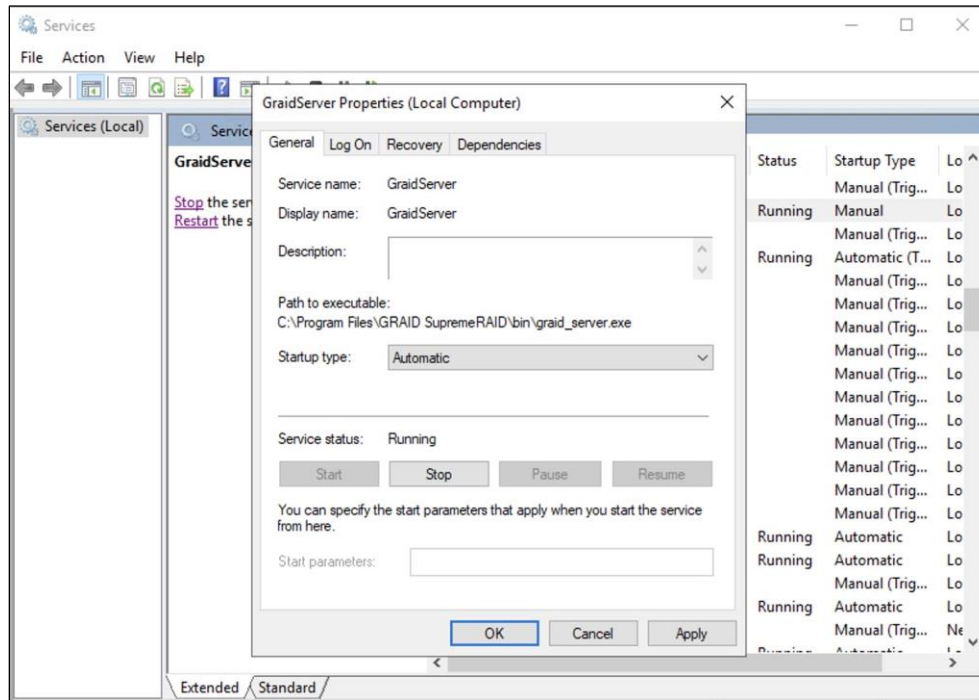
Output example: graid_server is not running.



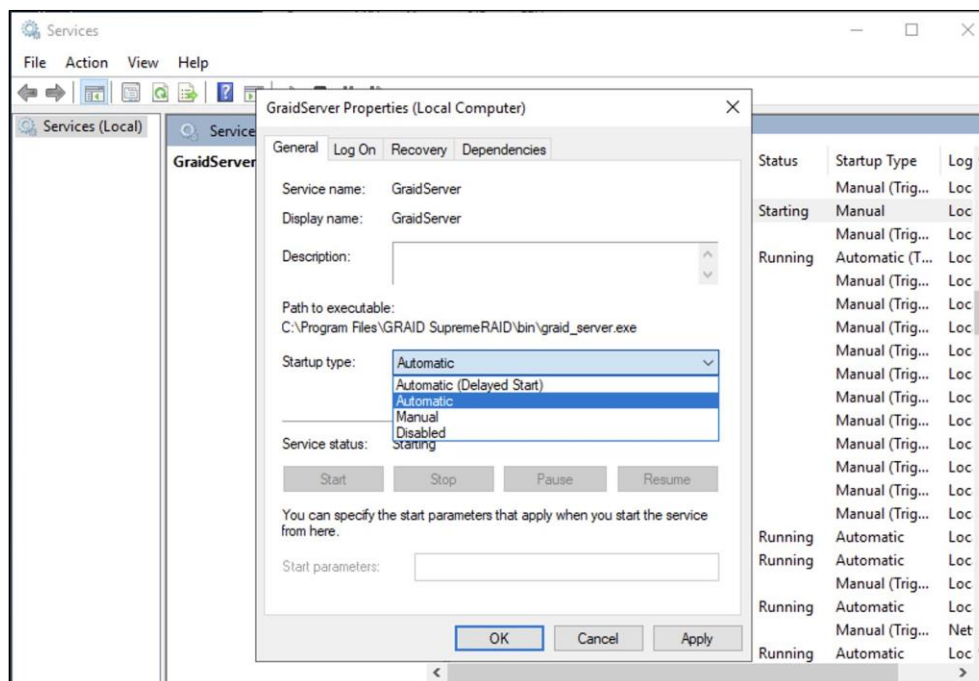
Starting graid_server:



graid_server is running:



Note: From the GraidServer Properties dialog, ensure that the Startup type: option is set to Automatic.



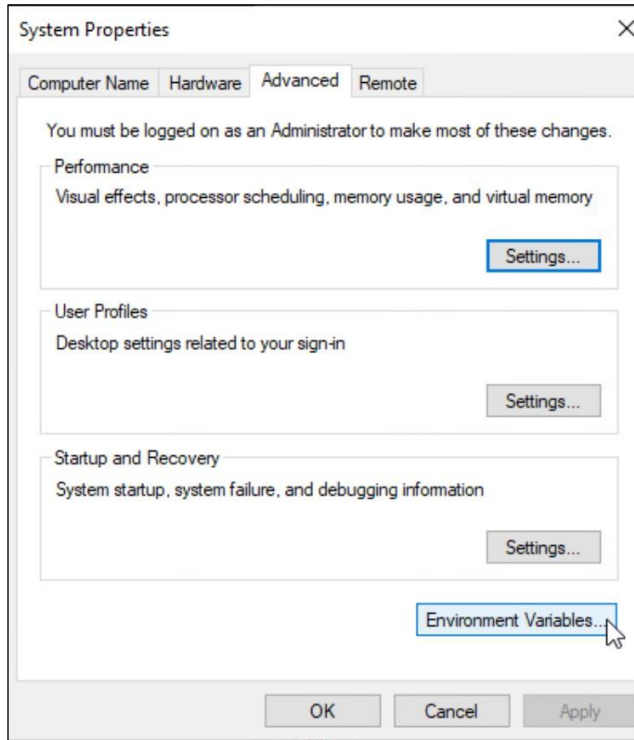
Different LED Blink Patterns on the Backplane

You might notice that the HDD/SSD activity indicator blink pattern is different on SupremeRAID™ than on traditional RAID cards.

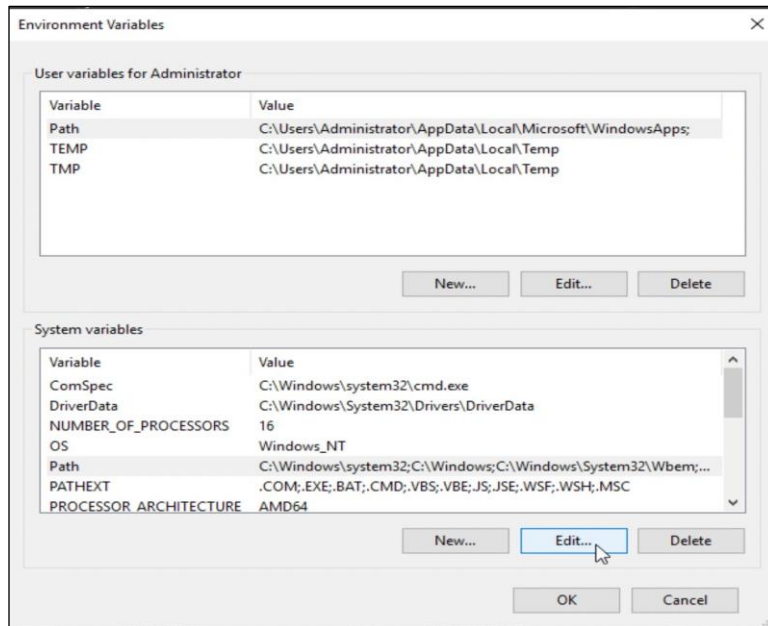
SupremeRAID™ does not require a buffering or caching mechanism to improve read/write performance like traditional RAID cards. This feature causes SupremeRAID™ indicators to blink differently than traditional RAID cards.

'graidctl' is not recognized as an internal or external command

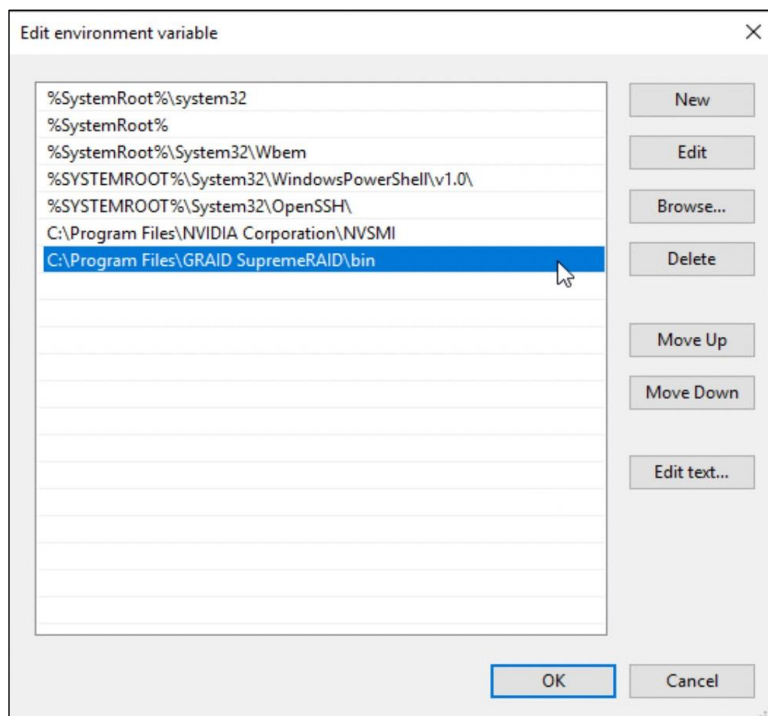
- A Press "Windows Key + R" and type "sysdm.cpl" to open the System Properties dialog.
- B Select the Advanced tab and click Environment Variables....



- C From the Environment Variables dialog box, select the **Path** row and click **Edit**.



D Add the SupremeRAID™ driver path, and then save it and reboot the system.



To change directories to the SupremeRAID™ driver folder: From the PowerShell or command prompt running as administrator, change to the directory containing the SupremeRAID™ driver folder.

SPECIFICATIONS

SAFETY INFORMATION

English Version

CE Directives Declaration: NVIDIA Corporation hereby declares that this device complies with all material requirements and other relevant provisions of the 2014/30/EU and 2011/65/EU. A copy of the Declaration of Conformity may be obtained directly from NVIDIA GmbH(Bavaria Towers - Blue Tower, Einstein Strasse 172, D-81677 Munich, Germany)

NVIDIA products are designed to operate safely when installed and used according to the product instructions and general safety practices. The guidelines included in this document explain the potential risks associated with equipment operation and provide important safety practices designed to minimize these risks. By carefully following the information contained in this document, you can protect yourself from hazards and create a safer environment.

This product is designed and tested to meet IEC 60950-1 and IEC 62368-1 Safety Standards for Information Technology Equipment. This also covers the national implementations of IEC 70950-1/62368-1 based safety standards around the world e.g. UL 62368-1. These standards reduce the risk of injury from the following hazards:

- Electric shock: Hazardous voltage levels contained in parts of the product
- Fire: Overload, temperature, material flammability
- Energy: Circuits with high energy levels (240-volt amperes) or potential as burn hazards.
- Heat: Accessible parts of the product at high temperatures.
- Chemical: Chemical fumes and vapors
- Radiation: Noise, ionizing, laser, ultrasonic waves

This device complies with part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) This device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

This product, as well as its related consumables and spares, complies with the reduction in hazardous substances provisions of the "India E-waste (Management and Handling) Rule 2016". It does not contain lead, mercury, hexavalent chromium, polybrominated biphenyls or polybrominated diphenyl ethers in concentrations exceeding 0.1 weight % and 0.01 weight % for cadmium, except for where allowed pursuant to the exemptions set in Schedule 2 of the Rule.

Retain and follow all product safety and operating instructions.

Always refer to the documentation supplied with your equipment. Observe all warnings on the product and in the operating instructions found on the product's User Guide.



This is a recycling symbol indicating that the product/battery cannot be disposed of in the trash and must be recycled according to the regulations and/or ordinances of the local community.



Hot surface warning. Contact may cause burns. Allow to cool before servicing.

Chinese Version (SC)

NVIDIA 产品在设计时充分考虑到操作安全性，可根据产品说明和常规安全做法进行安全安装和使用。本文档中包含的准则解释了设备操作所涉及的风险，并提供了最大限度降低这些风险的重要安全做法。请详细阅读本文档中的信息并按要求操作，这样可保护您免遭受为显并创建一个更加安全的环境。

本产品按照信息技术设备安全标准 IEC 60950-1 和 IEC 62368-1 进行设计，并且经测试表明符合这些设备。此处所述标准也包括全球各国/地区实施的基于 IEC 60950-1/62368-1 的安全标准，例如 UL 62328-1。这些标准降低了因以下危险而受伤的风险：

- 电击：部分产品中包含的危险电压水平起火：超载、高温、可燃性材料
- 机械：锋利的边缘、活动部件、不稳定性
- 电源：高电压电路（240 伏安）或潜在的烧伤风险
- 高温：产品的可触及部分存在高温化学：化学烟雾和蒸气
- 辐射：噪音、电离、激光、超声波

请牢记并遵守所有产品安全和操作说明。请务必参考您的设备随附的说明文档。请注意产品上以及产品用户指南的操作说明中列示的所有警告。



这是一个通用的回收标志，表示产品/电池不能以丢弃的方式处置，必须按造本地社区的法规和/或条例回收。



警告！表面高温。接触可能导致灼伤。请再冷却后再使用。



产品中有害物质的名称及含量根据中国电器电子产品有害物质限制使用管理办法)

内存	0	0	0	0	0	0
结构件以及风扇	x	0	0	0	0	0
线材/连接器	x	0	0	0	0	0
焊接金属	0	0	0	0	0	0
助焊剂，锡膏，标签及耗材	0	0	0	0	0	0
<p>本表格依据SJ/T 11364-2014的规定编制</p> <p>0：表示该有害物质在该部件所有的均质材料中的含量均在GB/T 26572 标准规定的限量要求以下。</p> <p>X：表示该有害物质至少在该部件的某一均质材料中的含量超出GB/T 26572标准规定的限量要求。</p> <p>此表中所有名称中含"x" 的部件均符合RoHS立法。</p> <p>注：环保使用期限的参考标识取决于产品正常工作的温度和湿度等条件</p>						

Chinese Version (TC)

在遵照產品說明與一般安全做法進行安裝與使用產品的情況下，NVIDIA 產品可安全地操作。本文件所列的準則說明與設備操作相關的潛在風險，同時也提供將這些風險降到最低的重要安全做法。謹慎遵守本文件中的資訊，您就可以避免危險並創造更安全的環境。

此產品係根據 Safety Standards for Information Technology Equipment(資訊技術設備安全標準) IEC 60950-1 和 IEC 62368-1 進行設計與測試。同時也涵蓋全世界國家以 IEC 60950-1/62368-1 為根據的安全標準，例如 UL 62368-1。這些標準可降低下列危險造成的傷害的風險：

- 觸電危險：本產品部分零件的電壓等級具危險性
- 火災危險：超載、溫度、材料可燃性
- 機械危險：尖銳邊緣、移動零件、不穩定性
- 電燒力危險：電路電壓高（240 電壓）或具有潛在起火燃燒熱能危險：產品表面可能達到高溫，注意燙傷危機
- 化學危險：化學異味氣體與蒸氣
- 輻射危險：噪音、游離輻射、雷射、超音波

請保留並遵守所有產品安全與操作說明的相關規定。請務必參閱設備隨附的文件。請遵守產品上，和產品使用者只能中操作說明裡的警告規定。



此國際回收標誌表示此產品/電池不能棄置於垃圾桶中，必須根據當地社區的規範和/或法令回收。



表面高溫警告。接觸時可能燙傷。使用前請先降溫。

限用物質含有情況標示聲明書						
設備名稱：繪圖卡						
單元	限用物質及其化學符號					
	鉛	汞	鎘	六價鉻	多溴聯苯	多溴二苯醚
PCB板	0	0	0	0	0	0
結構件以及風扇	-	0	0	0	0	0
連結器	-	0	0	0	0	0
被動電子零件	-	0	0	0	0	0
主動電子零件	-	0	0	0	0	0
內存	0	0	0	0	0	0
線材	0	0	0	0	0	0
焊接金屬	0	0	0	0	0	0
助焊劑、錫膏、標籤及耗材	0	0	0	0	0	0
備考1：0：係指該限用物質未超出百分比含量基準值。						
備考2：-：係指該限用物質為排除項目。						
此表中所有名稱含“-”的部件均符合歐盟RoHS立法。						
注：環保使用期限的參考標識取決於產品正常工作的溫度和濕度等條件。						